








A Unified Model for Real-Time Systems: Symbolic Techniques and Implementation

S. Akshay¹ , Paul Gastin^{2,4} , R. Govind¹ , Aniruddha R. Joshi¹ ,
and B. Srivathsan^{3,4} 

¹ Department of CSE, Indian Institute of Technology Bombay, Mumbai, India
{akshayss,govindr,aniruddhajoshi}@cse.iitb.ac.in

² Université Paris-Saclay, ENS Paris-Saclay, CNRS, LMF, 91190 Gif-sur-Yvette,
France

paul.gastin@ens-paris-saclay.fr

³ Chennai Mathematical Institute, Chennai, India
sri@cmi.ac.in

⁴ CNRS, ReLaX, IRL 2000, Siruseri, India



Abstract. In this paper, we consider a model of *generalized timed automata* (GTA) with two kinds of clocks, *history* and *future*, that can express many timed features succinctly, including timed automata, event-clock automata with and without diagonal constraints, and automata with timers.

Our main contribution is a new simulation-based zone algorithm for checking reachability in this unified model. While such algorithms are known to exist for timed automata, and have recently been shown for event-clock automata without diagonal constraints, this is the first result that can handle event-clock automata with diagonal constraints and automata with timers. We also provide a prototype implementation for our model and show experimental results on several benchmarks. To the best of our knowledge, this is the first effective implementation not just for our unified model, but even just for automata with timers or for event-clock automata (with predicting clocks) without going through a costly translation via timed automata. Last but not least, beyond being interesting in their own right, generalized timed automata can be used for model-checking event-clock specifications over timed automata models.

Keywords: Real-time systems · Timed automata · Event-clock automata · Clocks · Timers · Verification · Zones · Simulations · Reachability

This work was supported by UMI ReLaX, IRL 2000 and DST/CEFIPRA/INRIA Project EQuaVE. S Akshay was supported in part by DST/SERB Matrics Grant MTR/2018/000744. Paul Gastin was partially supported by ANR project Ticktac (ANR-18-CE40-0015).




Fig. 1. An automaton with clocks on left, and timers on right for same constraints.

1 Introduction

The idea of adding real-time dynamics to formal verification models started as a hot topic of research in the 1980s [6, 11]. Over the years, timed automata [8, 9] has emerged as a leading model for finite-state concurrent systems with real-time constraints. Timed automata make use of *clocks*, real-valued variables which increase along with time. Constraints over clock values can be used as guards for transitions, and clocks can be reset to 0 along transitions. It is notable that the early works in this area made use of *timers* to deal with real-time [13, 22, 32]. Timers are started by setting them to some initial value within a given interval. Their values decrease with time, and an *timeout* event can be used in transitions to detect the instant when the timers become 0. Quoting from [6], the shift from timers to clocks in timed automata, as we know them today, is attributed to the fact that: “*apart from some technical conveniences in developing the emptiness algorithm and proving its correctness, the reformulation allows a simple syntactic characterization of determinism for timed automata*”. Over the last thirty years, the study of timed automata has led to the development of rich theory and industry-strength verification tools. The use of clocks has also allowed for the extension of the model to more complex constraints and assignments to clocks in transitions [14, 17]. Furthermore, considering more sophisticated rates of evolution for clocks gives the yet another well-established model of hybrid automata [7].

When it comes to the reachability problem, timers do have some nice properties. Let us explain with an example. Figure 1 shows a timed automaton on the left, and an automaton with timers on the right, for the set of words ab^* such that the time between every consecutive letters is 1. The timed automaton sets clock x to 0 and checks for the guard $x = 1?$ to enforce the timing constraint. The automaton with timers, on the right, sets a timer t_x to 1, and asks for its expiry in the immediate next action. Clock y and timer t_y are not necessary for the required timing property, but we add them to illustrate a different aspect that we will describe now. To solve the reachability problem, a symbolic enumeration of the state space is performed. In the timed automaton, at state q_1 , the enumeration gives constraints $y - x = n$ for every $n \geq 0$. Starting from $y - x = n$ and executing b gives $y - x = n + 1$, due to the combination of guard $x = 1?$ and reset $x := 0$. This shows that a naïve symbolic enumeration is not bound

to terminate. The question of developing finite abstractions for timed automata has been a central problem of study which started in the late 90s and continues till date (see recent surveys [18,38]). Such an issue does not occur with timers. In the automaton with timers on the right, t_x is set to 1 and t_y is set to some arbitrary value in the transition to q_1 . This gives $-1 \leq t_y - t_x \leq \infty$ for the set of all possible timer values. When t_x times out, the value of t_y could still be any value from 0 to ∞ . When t_x is set to 1 again, the set of possible timer values still satisfies the same constraint $-1 \leq t_y - t_x \leq \infty$ leading to a fixed point with a finite reachable state space. The fact that symbolic enumeration terminates on an automaton with timers was already observed in [22]. To our knowledge, later works on timed automata reachability never went back to timers, and there is no tool support that we know of to deal with models with timers directly. We find this surprising given that timers occur naturally while modeling real-time systems and moreover they enjoy this finiteness property.

In addition to clocks and timers, *event-clocks* are another special type of clock variables that are used to deal with timing constraints [10], which are attached to events. An event-recording clock for event a maintains the time since the previous occurrence of a , whereas an event-predicting clock for a gives the time to the next occurrence of a . Event-clocks have been used in the model of event-clock automata (ECA), and also in the logic of event-clocks [36]. These works argue that event-clocks can express typical real-time requirements. Theoretically, ECA can be determinized, and hence complemented. Therefore, model-checking an event-clock (logic or automaton) specification φ over a timed automaton \mathcal{A} can be reduced to reachability on the product of \mathcal{A} and the ECA for $\neg\varphi$. This makes event-clocks a convenient feature in specifications.

Recently, a symbolic enumeration algorithm for ECA was proposed [3]. It was noticed that when restricted to event-predicting clocks, the symbolic enumeration terminates without any additional checks (similar to the case of timers), whereas for the combination involving event-recording clocks, one needs simulation techniques from the timed automata literature. The same work showed how to adapt the best known simulation technique from timed automata into the setting of ECA. However, as discussed above, for model-checking we need a model containing both conventional clocks, timers and event-clocks. To our knowledge, no tool can directly work on such models.

Our goal in this work is to provide a one stop solution to real-time verification, be it reachability analysis or model-checking (over event-clock specifications), be it using models with clocks, or models with timers. We consider a unified model of a timed automaton over variables that can simulate normal clocks, timers and event-clocks. Here are our key contributions:

1. We define a new model of generalized timed automata (GTA) which have two types of variables, called *history* clocks and *future* clocks. History clocks generalize normal clocks as well as event-recording clocks, while future clocks generalize event-predicting clocks and timers. However, unlike event-clocks, clocks in GTA are not necessarily associated with events. We also consider a generic syntax that allows for diagonal constraints between variables.

2. We show undecidability of reachability for GTA, and study a *safe subclass* that makes the model decidable. Safe GTA already subsume timed automata, event-clock automata (with diagonal constraints) and automata with timers.
3. We adapt state-of-the-art symbolic enumeration techniques from timed automata literature to safe GTA. While we make use of ideas presented in [22] and [3], these works do not contain diagonal constraints between variables. Our main technical and theoretical innovation lies in a new termination analysis of the symbolic enumeration in the presence of diagonal constraints. Surprisingly, we show that the enumeration terminates as long as the diagonal constraints are restricted to usual clocks and event-clocks, but not timers.
4. We develop a prototype implementation of our model and algorithm in TCHECKER, an open-source platform for timed automata analysis, and show promising results on several existing and new benchmarks. To the best of our knowledge, our tool is the first that can handle event-clock automata, a model that till date has been the subject of many theoretical results.

Related Works. In the work that first introduced ECA, a translation from ECA to a timed automaton was also proposed. However, this translation is not efficient: in the worst case, this translation incurs a blowup in the number of clocks and states. In [27, 28], an extrapolation approach using maximal constants has been studied for ECA. However, it has been observed that simulation-based techniques are both more effective [14, 16] and efficient [5, 24–26] than extrapolation for checking reachability. Recently, [3] proposed a zone-based reachability algorithm for diagonal-free ECA, using simulations for finiteness, but there was no accompanying implementation. Diagonal constraints have long been known to allow succinct modeling [15] for the class of timed-automata, but only recently a zone-based algorithm that directly works on such automata, was proposed. ECA with diagonals are more expressive than ECA [19]. In this work, we propose a zone-based algorithm for a unified model that subsumes ECA with diagonals.

The use of history clocks and prophecy clocks in ECAs is in the same spirit as past and future modalities in temporal logics - this makes ECAs an attractive model for writing timed specifications. Indeed, this has also led to a development of various temporal logics with event-clocks [1, 23, 36]. ECA with diagonal constraints have been well-studied, such as in the context of timeline based planning [19, 20]. Finally, while there has been substantial advances in the theory of ECA, to the best of our knowledge, the only tool that handles ECA is TEMPO [37], and even this tool is restricted to just history clocks.

Structure of the Paper. In Sect. 2 we start by defining the generalized model. Section 3 examines its expressiveness, while Sect. 4 deals with the reachability problem and the safe subclass. Section 5 develops the symbolic enumeration technique, while Sect. 6 explains how distance graphs can be extended to this setting. Section 7 is dedicated to finiteness. Finally, we provide our experimental results in Sect. 8 and conclude with Sect. 9. All the missing proofs can be found in the full version of the paper [2].

2 Generalized timed automata

In this section we introduce the unified model. While we build on classical ideas from timed automata, almost every aspect is extended and below we highlight these changes. We define $X = X_H \uplus X_F$ to be a finite set of real-valued variables called *clocks*, where X_H is the set of *history clocks*, and X_F is the set of *future clocks*. History clocks always have a non-negative value and can increase arbitrarily along with time. Future clocks always have a non-positive value and can only increase until their values hit 0. History clocks simulate the usual clocks in timed automata and recording clocks of event-clock automata (ECA), and future clocks simulate timers and prophecy clocks of ECA. Both these clocks can take a special “undefined value” which marks that they are inactive. To deal with this naturally, we consider an extension of the reals with $+\infty$ and $-\infty$ as in [3]. The difference here is that we also have the so-called diagonal constraints.

Extending Clock Constraints. Let $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, +\infty\}$ denote the set of all real numbers along with $-\infty$ and $+\infty$. The usual $<$ order on reals is extended to deal with $\{-\infty, +\infty\}$ as: $-\infty < c < +\infty$ for all $c \in \mathbb{R}$ and $-\infty < \infty$. Similarly, $\overline{\mathbb{Z}} = \mathbb{Z} \cup \{-\infty, +\infty\}$ denotes the set of all integers along with $-\infty$ and $+\infty$. Let $\mathbb{R}_{\geq 0}$ (resp. $\mathbb{R}_{\leq 0}$) be the set of non-negative (resp. non-positive) reals. Let $\mathcal{C} = \{(\triangleleft, c) \mid c \in \overline{\mathbb{R}} \text{ and } \triangleleft \in \{\leq, <\}\}$, called the set of weights.

Let $X \cup \{0\}$ be the set obtained by extending the clocks of GTA with the special constant clock 0. Note that this clock will always have the value 0. Let $\Phi(X)$ denote a set of clock constraints generated by the following grammar: $\varphi ::= x - y \triangleleft c \mid \varphi \wedge \varphi$ where $x, y \in X \cup \{0\}$, $(\triangleleft, c) \in \mathcal{C}$ and $c \in \overline{\mathbb{Z}}$. The introduction of the special constant clock 0 allows us to treat constraints with just a single clock as special cases: the constraint $x \triangleleft c$ is equivalent to $x - 0 \triangleleft c$ and the constraint $c \triangleleft x$ is equivalent to $0 - x \triangleleft -c$. We often write $x = c$ as a shorthand for $x \leq c \wedge c \leq x$. Constraints of the form $x - y \triangleleft c$ will be called *atomic constraints*. A constraint of the form $x - y \triangleleft c$ is a *diagonal* (resp. *non-diagonal*) constraint if $x, y \neq 0$ (resp. $x = 0$ or $y = 0$).

To evaluate the constraints allowed by $\Phi(X)$, we extend addition on real numbers with the convention that $(+\infty) + \alpha = \alpha + (+\infty) = +\infty$ for all $\alpha \in \overline{\mathbb{R}}$ and $(-\infty) + \beta = \beta + (-\infty) = -\infty$, as long as $\beta \neq +\infty$. We also extend the unary minus operation from real numbers to $\overline{\mathbb{R}}$ by setting $-(+\infty) = -\infty$ and $-(-\infty) = +\infty$. Abusing notation, we write $\beta - \alpha$ for $\beta + (-\alpha)$. Notice that with this extended addition, the minus operation does not distribute over addition¹.

Extending Valuations. A valuation of clocks is a function $v: X \cup \{0\} \mapsto \overline{\mathbb{R}}$ which maps the special clock 0 to 0, history clocks to $\mathbb{R}_{\geq 0} \cup \{+\infty\}$ and future clocks to $\mathbb{R}_{\leq 0} \cup \{-\infty\}$. We denote by $\mathbb{V}(X)$ or simply by \mathbb{V} the set of valuations over X . We say that clock x is *defined* (resp. *undefined*) in v when $v(x) \in \mathbb{R}$ (resp. $v(x) \in \{-\infty, +\infty\}$). Let $x, y \in X \cup \{0\}$ be clocks (including 0) and let (\triangleleft, c) be a weight. For valuations $v \in \mathbb{V}$, define $v \models y - x \triangleleft c$ as $v(y) - v(x) \triangleleft c$.

¹ Notice that $-(a + b) = (-a) + (-b)$ when a or b is finite or when $a = b$. But, when $a = +\infty$ and $b = -\infty$ then $-(a + b) = -\infty$ whereas $(-a) + (-b) = +\infty$.

We say that a valuation v satisfies a constraint φ in $\Phi(X)$, denoted as $v \models \varphi$, when v satisfies all atomic constraints in φ .

By definition, we easily check that the constraint $y - x \triangleleft c$ is equivalent to *true* (resp. *false*) when $(\triangleleft, c) = (\leq, +\infty)$ (resp. $(\triangleleft, c) = (<, -\infty)$). Constraints that are equivalent to *true* or *false* will be called *trivial*, whereas all others are *non-trivial* constraints. If $(\triangleleft, c) \neq (\leq, +\infty)$ then $v \models y - x \triangleleft c$ never holds when $v(x) = -\infty$. Also, if $v(x) = v(y) \in \{-\infty, +\infty\}$ then $v \models y - x \triangleleft c$ only holds for $(\triangleleft, c) = (\leq, +\infty)$. For a non-trivial constraint $y - x \triangleleft c$, we have

- $v \models y - x \triangleleft c$ iff $v(y) < +\infty = v(x)$ or $(v(x)$ is finite and $v(y) \triangleleft v(x) + c)$.
- $v \models y - x \leq -\infty$ iff $v(y) < +\infty = v(x)$ or $v(y) = -\infty < v(x)$.
- $v \models y - x < +\infty$ iff $v(x) \neq -\infty$ and $v(y) \neq +\infty$.

We abuse notation and for $Y \subseteq X$, we define $Y \triangleleft c$ as $\bigwedge_{y \in Y} y \triangleleft c$, and $Y = c$ as $\bigwedge_{y \in Y} y = c$. We denote by $v + \delta$ the *valuation* obtained from valuation v by increasing by $\delta \in \mathbb{R}_{\geq 0}$ the value of all clocks in X . Note that, from a given valuation, not all time elapse result in valuations since future clocks need to stay at most 0. For example, from a valuation with $v(x) = -3$ and $v(y) = -2$, where x, y are future clocks, one can elapse at most 2 time units.

Extending Resets. For history clocks, the reset operation sets the clock to 0. For future clocks, the reset operation says that all constraints on the clock must be discarded, i.e., the clock is *released*. Given that the set of clocks is partitioned into history clocks and future clocks, we use the same notation $[R]v$ to talk about the change of clocks in R , whether it be reset/release. Formally, given a set of clocks $R \subseteq X$, we define $[R]v$ as $\{v' \in \mathbb{V} \mid v'(x) = 0 \ \forall x \in R \cap X_H \text{ and } v'(x) = v(x) \ \forall x \notin R\}$. Observe that *the release operation* is implicit: each future clock in R could take any value (not necessarily the same) from $[-\infty, 0]$ in $[R]v$. Note that $[R]v$ is a singleton when R contains only history clocks - this corresponds exactly to the reset operation in timed automata. Then, we simply write $v' = [R]v$ instead of $\{v'\} = [R]v$. When R contains only future clocks, $[R]v$ is the set of valuations obtained by releasing each clock in R while keeping the value of all other clocks unchanged. For $W \subseteq \mathbb{V}$, we let $[R]W = \bigcup_{v \in W} [R]v$. We have $[R' \cup R'']W = [R']([R'']W)$.

Extending Guards and Transitions. Before we define GTA, let us focus on the language to specify transitions. In normal timed automata, as shown in Fig. 2, a transition reads a letter, checks a guard $g \in \Phi(X_H)$ and then resets a subset R of (history) clocks. But in any one transition only a pair of guard, reset is performed and one cannot interleave them.



Fig. 2. A transition of TA (left) and of a GTA (right)

We generalize this to our setting with history and future clocks but also to allow arbitrary interleaving of guards and changes (to model this with a TA one may use a sequence of multiple transitions without delays in-between.) Formally, an *instantaneous timed program* is generated by the following grammar:

$$\text{prog} := \text{guard} \mid \text{change} \mid \text{prog}; \text{prog}$$

where $\text{guard} = g \in \Phi(X)$ and $\text{change} = [R]$ for some $R \subseteq X$. While guard and change are atomic programs, $\text{prog}; \text{prog}$ refers to sequential composition. The set of all programs generated by the above grammar will be denoted Programs . Then on a transition, we simply have a pair of letter label and an instantaneous timed program, e.g., (a, prog) in Fig. 2 (right).

The semantics for programs on a transition must generalize semantics for guards (defined using satisfaction relation \models above) and resets/release (defined using $[R]$ above). But there is an obvious difference between these two: a guard may be crossed only if the valuation before the guard satisfies it, whereas a *change* (reset or release) defines a relation between the valuations before and after the change. To capture both in a uniform way, we define the semantics of programs as relations on pairs of valuations. Formally, for $v, v' \in \mathbb{V}$, $\text{prog} \in \text{Programs}$ we define $(v, v') \models \text{prog}$, more conveniently written as $v \xrightarrow{\text{prog}} v'$, inductively:

- $v \xrightarrow{g} v'$ if $v \models g$ and $v' = v$,
- $v \xrightarrow{[R]} v'$ if $v' \in [R]v$,
- $v \xrightarrow{\text{prog}_1; \text{prog}_2} v'$ if $\exists v'' \in \mathbb{V}$ such that $v \xrightarrow{\text{prog}_1} v''$ and $v'' \xrightarrow{\text{prog}_2} v'$.

Now, we have all the pieces necessary to define our generalized model.

Definition 1 (Generalized timed automata). A generalized timed automata \mathcal{A} is given by a tuple $(Q, \Sigma, X, \Delta, (q_0, g_0), (Q_f, g_f))$, where Q is a finite set of states, Σ is a finite alphabet of actions, $X = X_F \uplus X_H$ is a set of clocks partitioned into future and history clocks, the initialization condition is a pair comprising of an initial state $q_0 \in Q$ and an initial guard $g_0 \in \Phi(X)$ which should be satisfied by initial valuations, similarly, the final condition is a pair comprising of a set of final states $Q_f \subseteq Q$ along with a final guard g_f that must be satisfied by final valuations, and $\Delta \subseteq (Q \times \Sigma \times \text{Programs} \times Q)$ is a finite set of transitions. Δ contains transitions of the form (q, a, prog, q') , where q is the source state, q' is the target state, a is the action triggering the transition, and prog is the instantaneous timed program that is executed in sequence (from left to right) while firing the transition.

The semantics of a GTA $\mathcal{A} = (Q, \Sigma, X, \Delta, (q_0, g_0), (Q_f, g_f))$ is given by a transition system $\text{TS}_{\mathcal{A}}$ whose states are configurations (q, v) of \mathcal{A} , where $q \in Q$ and $v \in \mathbb{V}$ is a valuation. A configuration (q, v) is initial if $q = q_0$ and $v \models g_0$. A configuration (q, v) is accepting if $q \in Q_f$ and $v \models g_f$. Transitions of $\text{TS}_{\mathcal{A}}$ are of two forms: (1) *delay transition*: $(q, v) \xrightarrow{\delta} (q, v + \delta)$ if $(v + \delta) \models X_F \leq 0$, and

(2) *discrete transition*: $(q, v) \xrightarrow{t} (q', v')$ if $t = (q, a, \text{prog}, q') \in \Delta$ and $v \xrightarrow{\text{prog}} v'$. Thus, a discrete transition $t = (q, a, \text{prog}, q')$, where $\text{prog} = \text{prog}_1; \dots; \text{prog}_n$ can be taken from (q, v) if there are valuations v_1, \dots, v_n such that $v \xrightarrow{\text{prog}_1} v_1 \xrightarrow{\text{prog}_2} \dots \xrightarrow{\text{prog}_n} v_n = v'$. A *run* of a GTA is a finite sequence of transitions from an initial configuration of $\text{TS}_{\mathcal{A}}$. A run is said to be *accepting* if its last configuration is accepting.

3 Expressivity of GTA and Examples

The GTA model defined above is rather expressive. Figure 3 illustrates an example which accepts words of the form $a^n b^m$ with $m \leq n$, where each a occurs at time 0, after which b 's are seen one by one, with distance 1 between them. The history clock x is used to ensure the timing constraint. For every a that is read, the future clocks y, z decrease by 1. Hence the future clocks y, z maintain the opposite of the number of a 's seen. When the automaton starts reading b , the future clocks also start elapsing time and since they cannot go above 0, the number of b 's is at most the number of a 's. Such a language cannot be accepted by timed automata since the untimed language obtained by removing the time stamps needs to be regular in the case of timed automata. The GTA model is not only expressive, it is also convenient for use. To see this we now show that three classical models of timed systems can be easily captured using GTA. We also illustrate the modeling convenience provided by GTA in Sect. 8 based on experiments.




Fig. 3. Example of a GTA

Timed automata. Timed automata (TA) of Alur-Dill [9] can be modeled as a GTA as follows: (1) The set of states of the GTA is the same as the set of states of the TA. (2) There are no future clocks in the GTA and its history clocks are the clocks of the TA. (3) Each transition of the form $q \xrightarrow{a, g, R} q'$ in a TA, where g is a guard, a a letter and R a subset of clocks to be reset, is replaced by a transition $q \xrightarrow{a, \text{prog}} q'$ where $\text{prog} = \langle g; [R] \rangle$. (4) Initially, all clocks must be 0, captured by setting $g_0 = (X_H = 0)$. (5) The final guard is empty: $g_f = \text{True}$.

Event-clock Automata. Event-clock automata (ECA) of [10] can be modeled as a GTA as follows: (1) The set of states of the GTA is the same as the set of states of the ECA. (2) For each $a \in \Sigma$, the GTA has a history clock \overleftarrow{a} and a future clock \overrightarrow{a} . (3) Each transition of the form $q \xrightarrow{a, g} q'$ in a ECA, where

g is a guard of the ECA, a a letter, is replaced by a transition $q \xrightarrow{a, \text{prog}} q'$ where $\text{prog} := \langle (\vec{a} = 0); [\vec{a}]; g; [\overleftarrow{a}] \rangle$. (4) At initialization, history clocks must be undefined (set to ∞), captured by $g_0 = (X_H = \infty)$. (5) At acceptance, all future clocks must be undefined, i.e., $g_f = (X_F = -\infty)$.

Automata with Timers. The third model we consider is that of automata with timers. Timers are timing constructs that are started/initialized with a certain time value at some point/event and *count down* to 0. They measure the time from when they were started till the timer hits 0, where the event of hitting 0 is called *timeout*. However, they can be stopped using a *stop* event at any intermediate point instead and in which case the timer must be freed for reuse later. Timers are a common construct in protocol specification, e.g., the ITU standard which uses timers rather than clocks [30] and Mealy machines with timers [31].

In our setting, a timer can be seen as a specific instance of a future clock. More precisely Automata with timers ($A_{\overline{\Sigma}}$) can be modeled as GTA as follows: (1) The set of states of the GTA is the same as the set of states of $A_{\overline{\Sigma}}$. (2) The future clocks of GTA are the timers of $A_{\overline{\Sigma}}$ and there are no history clocks. Initially, the timers are undefined, captured by $g_0 = (X_F = -\infty)$ and $g_f = \text{True}$. (4) A transition of $A_{\overline{\Sigma}}$ with action a from q to q' is encoded as $q \xrightarrow{a, \text{prog}} q'$ with:

- if the transition starts timer x with value $c \in \mathbb{R}_{\geq 0}$, then $\text{prog} = \langle x = -\infty; [x]; x = -c \rangle$.
- if the transition is guarded by $\text{timeout}(x)$, then $\text{prog} = \langle x = 0; [x]; x = -\infty \rangle$.
- if the transition stops timer x , then $\text{prog} = \langle [x]; x = -\infty \rangle$.

We note that the timer above differs from a prophecy-event-clock (of ECA) though both are future clocks. Prophecy-clocks are released only when the event is seen, so at that point the value of the prophecy-clock must be 0. On the other hand timers can be stopped and released even when their value is not 0. This subtle difference has a surprising impact when we allow diagonal guards.

4 The Reachability Problem for GTA

We are interested in the *reachability problem* for GTA: given a GTA \mathcal{A} , does it have an accepting run? For normal TA, the reachability problem is decidable and PSPACE complete as shown in [9]. This was shown using the so-called region abstraction, by proving the existence of a finite time-abstract bisimulation. However, this is not the case for GTA. As explained in the previous subsection, GTA capture ECA, and as shown in [27, 28], there exists ECA for which there is no finite time-abstract bisimulation. However, reachability is still decidable in the specific case of ECA, as again shown in [10]. We note that for ECA model of [27, 28] there are no diagonal constraints. In this case they show decidability via zone-extrapolation. In [3], another approach for decidability via zone simulations is shown. But again even in this model diagonal constraints are disallowed. Even more critically in GTA, we can capture timers and a priori

we can have diagonal constraints even among timers. So, the question we ask is whether reachability is still decidable for GTA. Surprisingly, the answer is no. The intuition is that with future clocks and diagonal constraints, we get the ability to count (cf. Fig. 3).

Theorem 2. *Reachability for GTA is undecidable.*

Proof. We reduce from counter machines. Given a counter machine, we will build a GTA with one future clock y_C for each counter C and one extra future clock z . The reduction uses diagonal constraints between z and the future clocks y_C .

Initially and after each transition, the value of the future clock z will be 0. Since a future clock has to be non-positive, time elapse is impossible. As an invariant, the value of the future clock y_C is the opposite of the value of counter C . The operations on counter C are encoded with the following programs: (1) $\mathbf{zero}_C = \langle y_C = 0 \rangle$ (2) $\mathbf{inc}_C = \langle [z]; z = y_C - 1; [y_C]; y_C = z; [z]; z = 0 \rangle$ (3) $\mathbf{dec}_C = \langle y_C \leq -1; [z]; z = y_C + 1; [y_C]; y_C = z; [z]; z = 0 \rangle$. In programs \mathbf{inc}_C and \mathbf{dec}_C , each release of a future clock is followed by a constraint which restricts the value non-deterministically chosen during the release. For instance, $[z]; z = y_C - 1$ is equivalent to $z := y_C - 1$. Hence, the overall effect of \mathbf{inc}_C is $y_C := y_C - 1$, maintaining all other clocks unchanged, including the invariant $z = 0$. \square

Given this negative result, what can we do? A careful observation of the proof tells us that it is the interplay between diagonal constraints and arbitrary releases of future clocks that leads to undecidability. More precisely, the encoding depends on the fact that clocks z and y_C which are used in diagonal constraints ($z = y_C - 1$, $z = y_C + 1$ and $y_C = z$) may have arbitrary values when they are released. This suggests a restricted subclass that we formalize next.

Definition 3 (Safe GTA). *Let $X_D \subseteq X_F$ be a subset of future clocks.*

A program $\mathit{prog} = \langle g_1; [R_1]; g_2; [R_2]; \dots; g_k; [R_k]; g_{k+1} \rangle$ is X_D -safe if

- *diagonal constraints between future clocks are restricted to clocks in X_D : if $x - y \triangleleft c$ with $x, y \in X_F$ occurs in some g_i then $x, y \in X_D$;*
- *clocks in X_D should be 0 or $-\infty$ before being released: if $x \in X_D \cap R_i$ then $x = 0$ or $x = -\infty$ occurs in g_i .*

A GTA \mathcal{A} is X_D -safe if it only uses X_D -safe programs on its transitions and the initial guard g_0 sets each history clock to either 0 or ∞ .

Observe that the three examples discussed in Sect. 3 are safe. Timed automata do not have future clocks so the condition is vacuously true. In ECA, event-predicting clocks are always checked for 0 before being released, hence they are safe as well with $X_D = X_F$. Automata with timers without diagonal constraints are also trivially safe with $X_D = \emptyset$. The importance of safety is the following theorem which is the center-piece of this article.

Theorem 4. *Reachability for X_D -safe GTA is decidable.*

We will establish this theorem by showing a finite, sound and complete zone based reachability algorithm for X_D -safe GTA. If the given GTA is not X_D -safe, then we lose proof of termination (unsurprisingly, since the problem is undecidable), but we still maintain soundness. Thus, even for such GTA when our algorithm does terminate it will give the correct answer.

5 Symbolic Enumeration

We adapt the \mathcal{G} -simulation framework presented in [26] for timed automata with diagonal constraints to GTA. Diagonal constraints offer succinct modeling [15], but are quite challenging to handle efficiently in zone-based algorithms, and have led to pitfalls in the past: [14] showed that the erstwhile algorithm based on zone-extrapolations that was implemented in tools is incorrect for models with diagonal constraints; moreover no extrapolation based method can work for automata with diagonal constraints. The simulation framework by-passes this impossibility result and is the state-of-the-art for timed automata with diagonal constraints. The framework was extended to event-clock automata without diagonal constraints in [3]. We show that the ideas from [26] and [3] can be suitably combined to give an effective procedure for safe GTAs. This extension to GTAs enables us to understand the mechanics of diagonal constraints in future clocks.

The algorithm based on the \mathcal{G} -simulation framework involves:

1. computation of a set of constraints at every state of the automaton by a *static analysis* of the model,
2. a symbolic enumeration using *zones* to compute the *zone graph*,
3. a *simulation relation* between zones to ensure termination of the enumeration.

We will next adapt the static analysis to the GTA setting. The algorithm for the zone graph computation and the implementation of the simulation relation over zones is taken off-the-shelf from [26] and [3], except for a minor adaptation to include diagonal constraints involving future clocks. What is absent, and requires a non-trivial analysis, is the proof of termination. Therefore, we will mainly focus on this aspect and devote Sect. 7 for the termination argument.

\mathcal{G} -Simulation and the Static Analysis for GTA. We fix a GTA $\mathcal{A} = (Q, \Sigma, X, T, (q_0, g_0), (Q_f, g_f))$ for this section. Our goal is to define a simulation relation on the semantics of \mathcal{A} , i.e., on $\text{TS}(\mathcal{A})$. In the subsequent sections we will lift this to zones and show its finiteness. A simulation relation on $\text{TS}(\mathcal{A})$ is a reflexive, transitive relation $(q, v) \preceq (q', v')$ relating configurations with the same control state and (1) for every $(q, v) \xrightarrow{\delta} (q, v + \delta)$, we have $(q, v') \xrightarrow{\delta} (q, v' + \delta)$ and $(q, v + \delta) \preceq (q, v' + \delta)$, (2) for every transition t , if $(q, v) \xrightarrow{t} (q_1, v_1)$ for some valuation v_1 , then $(q, v') \xrightarrow{t} (q_1, v'_1)$ for some valuation v'_1 with $(q_1, v_1) \preceq (q_1, v'_1)$.

For any set G of atomic constraints, we define a *preorder* \preceq_G on valuations:

$$v \preceq_G v' \quad \text{if } \forall \varphi \in G, \forall \delta \geq 0, \quad v + \delta \models \varphi \implies v' + \delta \models \varphi.$$

Notice that in the definition above, we *do not* restrict δ to those such that $v + \delta$ is a valuation: we may have $v(x) + \delta > 0$ for some $x \in X_F$. In usual timed automata, this question does not arise, as elapsing any δ from any given valuation always results in a valuation. But this is crucial for the proof of Theorem 5 below.

Intuitively, the preorder above is a simulation wrt the constraints in G even after time elapse. But we need this to also be a simulation wrt discrete transitions. To achieve this, the set of constraints G should depend on the available

discrete transitions. In fact, we define a map \mathcal{G} from states to set of constraints, in such a way that it captures the simulation wrt the discrete actions. In other words, our focus will be to choose state-dependent sets of constraints (given by the map \mathcal{G}) depending on \mathcal{A} such that the resulting preorder induces a simulation on $\text{TS}(\mathcal{A})$.

As a first step towards this, we define, for any set G of constraints and any program prog , a set of constraints $G' = \text{pre}(\text{prog}, G)$ such that, if $v \preceq_{G'} v'$ and $v \xrightarrow{\text{prog}} v_1$ then there exists $v' \xrightarrow{\text{prog}} v'_1$ such that $v_1 \preceq_G v'_1$. This set is defined inductively as follows (G is a set of atomic constraints, R is a set of clocks, g is an *arbitrary* constraint, $y - x \triangleleft c$ is an *atomic* constraint):

$$\begin{aligned} \text{pre}(\text{prog}_1; \text{prog}_2, G) &= \text{pre}(\text{prog}_1, \text{pre}(\text{prog}_2, G)) \\ \text{pre}(g, G) &= \text{split}(g) \cup G \\ \text{pre}([R], G) &= \bigcup_{\varphi \in G} \text{pre}([R], \{\varphi\}) \end{aligned} \quad \text{pre}([R], \{y - x \triangleleft c\}) = \begin{cases} \{y - x \triangleleft c\} & \text{if } x, y \notin R \\ \{y \triangleleft c\} & \text{if } x \in R, y \notin R \\ \{-x \triangleleft c\} & \text{if } x \notin R, y \in R \\ \emptyset & \text{if } x, y \in R \end{cases}$$

where $\text{split}(g)$ is the set of atomic constraints occurring in g .

Now, the choice of suitable G will be obtained by static analysis, on the lines of what was done for timed automata with diagonals [24–26], but adapted to our more powerful model. More precisely, we define the map \mathcal{G} from Q to sets of atomic constraints as the least fixpoint of the set of equations:

$$\mathcal{G}(q) = \{x \leq 0 \mid x \in X_F\} \cup \bigcup_{q \xrightarrow{a, \text{prog}} q'} \text{pre}(\text{prog}, \mathcal{G}(q')) \quad (1)$$

Finally, based on \preceq_G and the $\mathcal{G}(q)$ computation, we can define a preorder $\preceq_{\mathcal{A}}$ between configurations of $\text{TS}(\mathcal{A})$ as $(q, v) \preceq_{\mathcal{A}} (q', v')$ if $q = q'$ and $v \preceq_{\mathcal{G}(q)} v'$. We then show that $\preceq_{\mathcal{A}}$ defined above is indeed a simulation relation.

Theorem 5. *The relation $\preceq_{\mathcal{A}}$ is a simulation on the transition system $\text{TS}_{\mathcal{A}}$.*

Zones for GTA and the Zone Graph Computation. Roughly, zones [12] are sets of valuations that can be represented efficiently using constraints between differences of clocks. In this section, we introduce an analogous notion for generalized timed automata. We consider *GTA zones*, or simply *zones*, which are special sets of valuations of GTA. A GTA zone is a set of valuations satisfying a conjunction of constraints of the form $y - x \triangleleft c$, where $x, y \in X \cup \{0\}$, $c \in \overline{\mathbb{Z}}$ and $\triangleleft \in \{\leq, <\}$. Thus zones are an abstract representation of sets of valuations. Then, an abstract configuration, also called a *node*, is a pair consisting of a state and a zone. Firing a transition $t := (q, a, \text{prog}, q')$ in a GTA \mathcal{A} from node (q, Z) will result in another node following a sequence of operations that we now define. *GTA Zone Operations.* Let g be a guard, $R \subseteq X$ a set of clocks and Z a GTA zone.

- Guard intersection: $Z \cap g := \{v \mid v \in Z \text{ and } v \models g\}$
- Release/Reset: $[R]Z = \bigcup_{v \in Z} [R]v$ (as defined in Sect. 2)
- Time elapse: $\vec{Z} = \{v + \delta \mid v \in Z, \delta \in \mathbb{R}_{\geq 0} \text{ s.t. } v + \delta \models (X_F \leq 0)\}$

Successor Computation. We can show that starting from a zone Z , the successors after the above operations are also zones (see Theorem 29 in [2]). A guard g can be seen as yet another zone and hence guard intersection is just an intersection operation between two zones. Similarly, the change operation preserves zones. Finally, as is usual with timed automata, zones are closed under the time elapse operation.

Thus, for a transition $t := (q, a, \text{prog}, q')$ and a node (q, Z) , we can define the successor node (q', Z') , and we write $(q, Z) \xrightarrow{t} (q', Z')$, where Z' is the zone computed by the following sequence of operations: Let $\text{prog} = \text{prog}_1; \dots; \text{prog}_n$, where each prog_i is an atomic program, i.e., a guard or a change. Then we define zones Z_1, \dots, Z_{n+1} where, $Z_1 = Z$, $Z' = \overrightarrow{Z_{n+1}}$, and for each $1 \leq i \leq n$, $Z_{i+1} = Z_i \cap g_i$ if prog_i is a guard g_i , and $Z_{i+1} = [R_i]Z_i$ if prog_i is a change $[R_i]$.

Now, we can lift zone graphs, simulations from TA to GTA and obtain a symbolic reachability algorithm for GTA.

Definition 6 (GTA zone graph). *Given a GTA \mathcal{A} , its GTA zone graph, denoted $\text{GZG}(\mathcal{A})$, is defined as follows: Nodes are of the form (q, Z) where q is a state and Z is a GTA zone. The initial node is $(q_0, \overrightarrow{Z_0})$ where q_0 is the initial state and Z_0 is the set of all valuations which satisfy the initial constraint g_0 : Z_0 is given by $g_0 \wedge (X_F \leq 0) \wedge (X_H \geq 0)$. For every node (q, Z) and every transition $t := (q, a, \text{prog}, q')$ of \mathcal{A} , there is a transition $(q, Z) \xrightarrow{t} (q', Z')$ in the GTA zone graph. A node (q, Z) is accepting if $q \in Q_f$ and $Z \cap g_f$ is non-empty, i.e., there exists a valuation in Z satisfying the final constraint.*

Similar to the case of zone graphs for timed automata and event zone graphs for ECA, the GTA zone graph can be used to decide reachability for generalized timed automata. A node (q, Z) is said to be reachable (in \mathcal{A}) if there is a path from the initial node $(q_0, \overrightarrow{Z_0})$ to (q, Z) in $\text{GZG}(\mathcal{A})$. Thus, reachability of a final state in \mathcal{A} reduces to checking reachability of an accepting node in $\text{GZG}(\mathcal{A})$. However, as in the case of zone graphs for timed automata, $\text{GZG}(\mathcal{A})$ is also not guaranteed to be finite. Hence, we need to compute a finite truncation of the GTA zone graph, which is still sound and complete for reachability.

Definition 7 (Simulation on GTA zones and finiteness). *Let \preceq be a simulation relation on $\text{TS}(\mathcal{A})$. For two GTA zones Z, Z' , we say $(q, Z) \preceq (q, Z')$ if for every $v \in Z$ there exists $v' \in Z'$ such that $(q, v) \preceq (q, v')$. The simulation \preceq is said to be finite if for every sequence $(q, Z_1), (q, Z_2), \dots$ of reachable nodes, there exists $j > i$ such that $(q, Z_j) \preceq (q, Z_i)$.*

Now, the reachability algorithm, as in TA, enumerates the nodes of the GTA zone graph and uses the simulation $\preceq_{\mathcal{A}}$ from Theorem 5 to truncate nodes that are smaller with respect to the simulation. In Sect. 7, we will show that $\preceq_{\mathcal{A}}$ is finite when \mathcal{A} is safe, which implies that the reachability algorithm terminates. But before that we discuss the issue of implementability.

6 Computing with GTA Zones Using Distance Graphs

To implement the reachability algorithm described above, we will view zones as *distance graphs*, as is usually done in the literature [12].

Recall the notion of weights $\mathcal{C} = \{(\triangleleft, c) \mid c \in \overline{\mathbb{R}} \text{ and } \triangleleft \in \{\leq, <\}\}$. An order relation $<$ between weights is defined as $(\triangleleft, c) < (\triangleleft', c')$ when either (1) $c < c'$, or (2) $c = c'$ and \triangleleft is $<$ while \triangleleft' is \leq . Note that since $(<, -\infty) < (\leq, -\infty) < (\triangleleft, c) < (<, \infty) < (\leq, \infty)$ for all $c \in \mathbb{R}$, this relation is a total order and therefore \min of a finite set of weights is well defined. We also use the commutative and associative sum operation on weights defined in [4]. If $c, c' \in \mathbb{R}$ are finite, the definition is as usual: $(\triangleleft, c) + (\triangleleft', c') = (\triangleleft'', c + c')$ where $\triangleleft'' = \leq$ if $\triangleleft = \triangleleft' = \leq$ and $\triangleleft'' = <$ otherwise. Infinite weights α, β from the list $(<, +\infty), (\leq, -\infty), (\leq, +\infty), (<, -\infty)$ are all ‘absorbants’ wrt. weaker weights: $\alpha + \beta = \beta + \alpha = \alpha$ if α is stronger than β (i.e., α is listed after β). Also, $\alpha + (\triangleleft, c) = \alpha$ if $c \in \mathbb{R}$ is finite.

A distance graph \mathbb{G} is a weighted directed graph without self-loops, with vertex set $X \cup \{0\} = X_F \cup X_H \cup \{0\}$, and edges labeled with weights from $\mathcal{C} \setminus \{(<, -\infty)\}$. We define its semantics $\llbracket \mathbb{G} \rrbracket := \{v \in \mathbb{V} \mid v \models y - x \triangleleft c \text{ for all edges } x \xrightarrow{\triangleleft c} y \text{ in } \mathbb{G}\}$. The weight of edge $x \rightarrow y$ is denoted \mathbb{G}_{xy} and we set $\mathbb{G}_{xy} = (\leq, \infty)$ if there is no edge $x \rightarrow y$. The weight of a path is the sum of the weights of its edges. A cycle in \mathbb{G} is said to be negative if its weight is strictly less than $(\leq, 0)$.

In classical timed automata, the significance of distance graphs stems from the observation that a distance graph has no negative cycles iff its semantics is non-empty. This property does not immediately hold for distance graphs over the extended algebra [4, Section 4.2] However, we can convert a distance graph \mathbb{G} (in time polynomial in number of clocks) into a *standard form* where this characterization continues to hold. First, we set $\mathbb{G}'_{0x} = \min(\mathbb{G}_{0x}, (\leq, 0))$ for $x \in X_F$ and $\mathbb{G}'_{x0} = \min(\mathbb{G}_{x0}, (\leq, 0))$ for $x \in X_H$. Moreover, if $x \in X_F$ then we set $\mathbb{G}'_{x0} = \min(\mathbb{G}_{x0}, (<, \infty))$ if $\mathbb{G}_{xy} \neq (\leq, \infty)$ for some $y \neq x$, otherwise we keep $\mathbb{G}'_{x0} = \mathbb{G}_{x0}$. Similarly, if $y \in X_H$ then we set $\mathbb{G}'_{0y} = \min(\mathbb{G}_{0y}, (<, \infty))$ if $\mathbb{G}_{xy} \neq (\leq, \infty)$ for some $x \neq y$, otherwise we keep $\mathbb{G}'_{0y} = \mathbb{G}_{0y}$. Finally, for $x, y \in X$ with $x \neq y$ we set $\mathbb{G}'_{xy} = \mathbb{G}_{xy}$. The graph \mathbb{G}' constructed above is called the standardization of \mathbb{G} , it is equivalent to \mathbb{G} (i.e., $\llbracket \mathbb{G}' \rrbracket = \llbracket \mathbb{G} \rrbracket$) and it has a negative cycle iff its semantics $\llbracket \mathbb{G}' \rrbracket$ is empty [4].

Now, suppose \mathbb{G}' (in standard form) has no negative cycles, then we construct \mathbb{G}'' by replacing the weight of an edge $x \rightarrow y$ by the minimum of the weights of the paths from x to y in \mathbb{G}' . Such a \mathbb{G}'' is called the *normalization* of \mathbb{G}' and has several useful properties.

Let Z be a nonempty zone. Writing the constraints in Z as a distance graph, followed by standardizing and normalizing it, results in *its canonical distance graph* $\mathbb{G}(Z)$: $\llbracket \mathbb{G}(Z) \rrbracket = Z$ and $\mathbb{G}(Z)$ is minimal among the standard graphs G with $\llbracket G \rrbracket = Z$. We denote by Z_{xy} the weight of the edge $x \rightarrow y$ in $\mathbb{G}(Z)$.

[3] contains the algorithms for the zone operations when there are no diagonal constraints. Successor computation can be done in $\mathcal{O}(|X|^2 \cdot |g|)$ and the simulation in $\mathcal{O}(|X|^2)$. Incorporating intersection with diagonal constraints requires an additional standardization step since diagonal constraints may break this property. A detailed explanation of the successor computation of zones is provided in

[2]. For the simulation, the algorithm from [26] is used. However, in the presence of diagonal constraints, the simulation check becomes NP-complete in general, and makes use of heuristics that allows for a faster check in practice. What remains is to show that $\preceq_{\mathcal{A}}$ is a finite simulation for X_D -safe GTA.

7 Finiteness of the Simulation Relation

In this section, we show that the simulation relation $\preceq_{\mathcal{A}}$ proposed in Sect. 5 is finite for safe GTA, which proves termination of the symbolic enumeration-based reachability algorithm. We do this in two parts: first, we show that the zones that are reached during the enumeration satisfy some invariants, in particular, only finitely many values occur in constraints among future clocks. This is however not necessarily true for history clocks. There the simulation comes into play. In the second part of the proof, we combine the invariants with an equivalence relation to show finiteness of the simulation. Below, we sketch these arguments and provide intuition leaving formal details to [2] due to lack of space.

Throughout this section, we fix an X_D -safe GTA \mathcal{A} . Let $M = \max\{|c| \mid c \in \mathbb{Z} \text{ is used in some constraint of } \mathcal{A}\}$, called the maximal constant of \mathcal{A} . We say that a zone Z is reachable if there is some reachable node (q, Z) in $\text{GZG}(\mathcal{A})$.

Part 1: Invariants on zones. We start by showing an important property of reachable zones: closure under valuations that agree on the value of history clocks, and satisfy the same set of safe constraints involving non-history clocks.

We say that a constraint $x - y \triangleleft c$ is M -bounded if either $c \in \mathbb{R}$ is such that $|c| \leq M$ or $c \in \{-\infty; +\infty\}$. It is X_D -safe if $x, y \in X_F$ implies $x, y \in X_D$. We say that it is (X_D, M) -safe if it is both M -bounded and X_D -safe.

Lemma 8. *Let $v, v' \in \mathbb{V}$ be such that $v' \downarrow_{X_H} = v \downarrow_{X_H}$ and, for all (X_D, M) -safe constraints $y - x \triangleleft c$ with $x, y \in X_F \cup \{0\}$, we have $v' \models y - x \triangleleft c$ if and only if $v \models y - x \triangleleft c$. Let Z be a reachable zone. Then, $v \in Z$ if and only if $v' \in Z$.*

The proof (given in [2]) works by establishing that the property is true in the initial zone, and showing that it is invariant under the zone operations used to compute $\text{GZG}(\mathcal{A})$. This proof crucially uses the fact that \mathcal{A} is X_D -safe. For the case of releasing a clock $x \in X_F \setminus X_D$, we use the fact that a diagonal constraint involving x may not use another future clock. For the case of releasing a clock $x \in X_D$, we use the fact that the value of the clock must be 0 or $-\infty$ just before the release. As a non-example, consider Fig. 3. Here, $X_D = \{y, z\}$ and $M = 1$. After two iterations of a , the zone Z_2 reached is $x = 0 \wedge y = z = -2$. Pick $v : x = 0, y = z = -2$ and $v' : x = 0, y = z = -3$. Notice that both of them satisfy the same set of (X_D, M) -safe constraints, but $v \in Z_2, v' \notin Z_2$. Indeed, the automaton is not X_D -safe since y and z are released arbitrarily.

From Lemma 8, we get the following corollary (with a more precise statement and proof in [2]). Namely, if a reachable zone Z contains a valuation v in which the difference between two future clocks x, y (including the zero clock) is finite and large enough, then Z contains valuations where the difference between x and y is any finite and large enough value.

Corollary 9. *Let Z be a reachable zone and let $v \in Z$. Let $n = \max(1, |X_D|)$. For all $x, y \in X_F \cup \{0\}$, if $-\infty < v(x) - v(y) < -nM$ then, for every α with $-\infty < \alpha < -nM$, we have a valuation $v' \in Z$ with $v'(x) - v'(y) = \alpha$.*

Notice that the property above does not hold if we simply take $n = 1$. For instance, if we have two clocks $x, z \in X_D$ then, applying the (X_D, M) -safe program $\langle [x, z]; z = -M \wedge x - z = -M \rangle$ from \mathbb{V} results in a zone Z where all valuations v satisfy $v(x) = -2M$. So the property fails with $n = 1$, x and $y = 0$. This is a noteworthy difference between models with and without diagonals.

Using Corollary 9, we can prove the main invariants satisfied by the zones obtained during the enumeration. Essentially, the weights of edges involving non-history clocks come from a finite set which depends on the number of future clocks in X_D and the maximum constant M of the automaton. This also induces an invariant on the constraint between a history clock and a future clock.

Before stating the result, we first give two technical lemmas from [4] that we use extensively in the proof.

Lemma 10 ([4]).

1. Let (\triangleleft, c) be a weight and $\alpha \in \overline{\mathbb{R}}$. Then,
 - $\alpha \triangleleft c$ iff $(\leq, \alpha) \leq (\triangleleft, c)$ iff $(\leq, 0) \leq (\leq, -\alpha) + (\triangleleft, c)$,
 - $\alpha \not\triangleleft c$ iff $(\triangleleft, c) < (\leq, \alpha)$ iff $(\leq, -\alpha) + (\triangleleft, c) < (\leq, 0)$ iff $(\leq, -\alpha) + (\triangleleft, c) \leq (<, 0)$.
2. Let $(\triangleleft, c), (\triangleleft', c'), (\triangleleft'', c'')$ be weights with $(\leq, 0) \leq (\triangleleft, c) + (\triangleleft', c')$. Then, there exists $\alpha \in \overline{\mathbb{R}}$ such that $\alpha \triangleleft c$ and $-\alpha \triangleleft' c'$. If in addition we have $(\triangleleft'', c'') < (\triangleleft, c)$ then there exists such an α with $\alpha \not\triangleleft'' c''$.

Lemma 11 ([4]). Let $\mathbb{G} = \mathbb{G}(Z)$ for a non-empty GTA zone Z , and let $x, y \in X \cup \{0\}$ be a pair of distinct nodes and $\alpha \in \overline{\mathbb{R}}$. There is a valuation $v \in \llbracket \mathbb{G} \rrbracket$ with $v(y) - v(x) = \alpha$ if and only if

1. $(\leq, \alpha) \leq \mathbb{G}_{xy}$ and $(\leq, -\alpha) \leq \mathbb{G}_{yx}$, and
2. if $x, y \in X$ and $\alpha \in \mathbb{R}$ is finite then the weights $\mathbb{G}_{x0}, \mathbb{G}_{0x}, \mathbb{G}_{y0}, \mathbb{G}_{0y}$ are all different from $(\leq, -\infty)$, and
3. if $x, y \in X$ and $\alpha = -\infty$ then $\mathbb{G}_{0x} \neq (\leq, -\infty) \neq \mathbb{G}_{y0}$.

Lemma 12. Let Z be a nonempty reachable zone. Let $n = \max(1, |X_D|)$. Then, the normalized distance graph $\mathbb{G}(Z)$ satisfies the following (\dagger) conditions:

- \dagger_1 For all $x \in X_F, y \in X_H \cup \{0\}$, if Z_{xy} is finite, then $(\leq, 0) \leq Z_{x0} \leq (\leq, nM)$.
- \dagger_2 For all $x \in X_F$, if Z_{0x} is finite, then $(<, -nM) \leq Z_{0x} \leq (\leq, 0)$.
- \dagger_3 For all $x \in X_H, y \in X_F$, if Z_{0y} is finite, then $Z_{x0} + (<, -nM) \leq Z_{xy}$.
- \dagger_4 For $x, y \in X_F$, if Z_{xy} is finite, then $(<, -nM) \leq Z_{xy} \leq (\leq, nM)$.

Proof. We focus on \dagger_1, \dagger_2 , leaving the more complicated cases to [2].

- \dagger_1 First, we consider the case where $y = 0$. So we assume that $(\leq, 0) \leq Z_{x0} < (<, \infty)$ is finite. Towards a contradiction, suppose that $(\leq, nM) < Z_{x0} < (<, \infty)$. Since Z is non-empty, we know that $(\leq, 0) \leq Z_{x0} + Z_{0x}$. Then, using

Lemma 10, we can find $\alpha \in \overline{\mathbb{R}}$ such that $(\leq, \alpha) \leq Z_{x0}$, $(\leq, -\alpha) \leq Z_{0x}$, and $nM < \alpha$. Notice that $\alpha < \infty$ since $Z_{x0} < (<, \infty)$. Further, using Lemma 11, we can get a valuation $v \in Z$ such that $0 - v(x) = \alpha$. Since $nM < \alpha < \infty$, this implies $-\infty < v(x) < -nM$. Let $Z_{x0} = (\triangleleft, c)$. We have $nM < c < \infty$. Using Corollary 9, we can get a valuation $v' \in Z$, such that $-\infty < v'(x) < -c$, a contradiction as it violates the constraint $0 - x \triangleleft c$ of Z . Next, assume that $Z_{xy} < (<, \infty)$ for some $y \in X_H$. Since Z is normal, we have $Z_{x0} \leq Z_{xy} + Z_{y0} < (<, \infty)$ as $Z_{xy} < (<, \infty)$ and $Z_{y0} \leq (\leq, 0)$. We now conclude from the first case that $(\leq, 0) \leq Z_{x0} \leq (\leq, nM)$.

†₂ We have to show that either $Z_{0x} = (\leq, -\infty)$ or $(<, -nM) \leq Z_{0x} \leq (\leq, 0)$. Let $Z_{0x} = (\triangleleft, c)$. Suppose $(\leq, -\infty) < Z_{0x} < (<, -nM)$. We have $-\infty < c < -nM$. As before, we can find α such that $(\leq, \alpha) \leq Z_{0x}$, $(\leq, -\alpha) \leq Z_{x0}$ and $\alpha \neq -\infty$. Then, by Lemma 11, we can find $v \in Z$ with $v(x) = \alpha$. We have $-\infty < v(x) \triangleleft c < -nM$. Now, using Corollary 9, we can get a valuation $v' \in Z$ such that $c < v'(x) < -nM$, which leads to a contradiction as it violates the constraint $x - 0 \triangleleft c$ in the zone. □

Part 2. Equivalence and Finiteness. We introduce below an equivalence relation \sim_M^n of *finite index* on valuations, depending on $n = \max(1, |X_D|)$ and the maximal constant M , and show that, if G is a set of atomic M -bounded integral constraints and if Z is a zone such that its canonical distance graph $\mathbb{G}(Z)$ satisfies (†) conditions, then the downward closure $\downarrow_G Z = \{v \in \mathbb{V} \mid \exists v' \in Z \text{ with } v \preceq_G v'\}$ is a union of \sim_M^n equivalence classes.

First, we define \sim_M on $\alpha, \beta \in \overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$ by $\alpha \sim_M \beta$ if $(\alpha \triangleleft c \iff \beta \triangleleft c)$ for all (\triangleleft, c) with $\triangleleft \in \{<, \leq\}$ and $c \in \{-\infty, \infty\} \cup \{d \in \mathbb{Z} \mid |d| \leq M\}$. In particular, if $\alpha \sim_M \beta$ then $(\alpha = -\infty \iff \beta = -\infty)$ and $(\alpha = \infty \iff \beta = \infty)$.

Next, for valuations $v_1, v_2 \in \mathbb{V}$, we define $v_1 \sim_M^n v_2$ by two conditions: $v_1(x) \sim_{nM} v_2(x)$ and $v_1(x) - v_1(y) \sim_{(n+1)M} v_2(x) - v_2(y)$ for all clocks $x, y \in X$. Notice that we use $(n + 1)M$ for differences of values. Clearly, \sim_M^n is an equivalence relation of finite index on valuations. Using this, we can show that the zones that are reachable in a safe GTA are unions of \sim_M^n -equivalence classes.

Lemma 13. *Let G be a set of X_D -safe M -bounded integral constraints which contains both $x \leq 0$ and $0 \leq x$ for each future clock $x \in X_F$. Let Z be a zone with a canonical distance graph $\mathbb{G}(Z)$ satisfying the (†) conditions of Lemma 12. Let $v_1, v_2 \in \mathbb{V}$ be valuations with $v_1 \sim_M^n v_2$. Then, $v_1 \in \downarrow_G Z$ iff $v_2 \in \downarrow_G Z$.*

Finally, from Lemmas 12 and 13, we obtain our main theorem of the section.

Theorem 14. *The simulation relation $\preceq_{\mathcal{A}}$ is finite if \mathcal{A} is safe.*

Proof. Let $(q, Z_0), (q, Z_1), (q, Z_2), \dots$ be an infinite sequence of *reachable* nodes in the zone graph of \mathcal{A} . By Lemma 12, for all i , the distance graph $\mathbb{G}(Z_i)$ in canonical form satisfies conditions (†).

The set $\mathcal{G}(q)$ contains only X_D -safe and M -bounded integral constraints. Let G be $\mathcal{G}(q)$ together with the constraints $x \leq 0$ and $0 \leq x$ for each future clock

Table 1. Experimental results obtained by running our prototype implementation and, when possible, the standard reachability algorithm using \mathcal{G} -simulation implemented in TCHECKER. Both implementations use a breadth-first search with simulation. For each model, we give the parameters in parenthesis - for ToyECA, we explain the parameterization in [2], while for others, we report the number of concurrent processes. All experiments were run on an Ubuntu machine with an Intel-i5 7th Generation processor and 8 GB RAM, and timeout set to 60s.

Sl. No.	Models	\mathcal{G} -Sim			GTA Reach		
		Visited nodes	Stored nodes	Time in sec	Visited nodes	Stored nodes	Time in sec
1	Dining Phi. (6)	5480	5480	4.911	5480	5480	6.410
2	FDDI (10)	10219	459	10.139	10219	459	16.797
3	Fischer (10)	447598	260998	29.1574	447598	260998	34.6517
4	ToyECA(10000, 4)	150049	49	4.22	3	3	0.0003
5	ToyECA(5000, 6)	315193	193	15.572	3	3	0.0006
6	ToyECA(1000, 100)	TIMEOUT			3	3	0.877
7	ToyECA(50000, 120)	TIMEOUT			3	3	1.52
8	Fire-alarm-pattern(5)	—			46	46	0.027
9	CSMACD-bounded(1)	—			34	26	0.0054
10	CSMACD-bounded(4)	—			4529	2068	2.597
11	ABP-prop1(1)	—			114	114	0.038
12	ABP-prop2(1)	—			168	168	0.026

$x \in X_F$. From Lemma 13 we deduce that for all i , $\downarrow_G Z_i$ is a union of \sim_M^n -classes. Since \sim_M^n is of finite index, there are only finitely many unions of \sim_M^n -classes. Therefore, we find $i < j$ with $\downarrow_G Z_i = \downarrow_G Z_j$, which implies $Z_j \preceq_G Z_i$. Since $\mathcal{G}(q) \subseteq G$, this also implies $Z_j \preceq_{\mathcal{G}(q)} Z_i$. \square

8 Experimental Evaluation

We have implemented a prototype that takes as input a GTA, as given in Definition 1, and applies our reachability algorithm, in the open source tool TCHECKER [29]. To do so, we extend TCHECKER to allow clocks to be declared as one of *normal*, *history*, *prophecy*, or *timer*, and extend the syntax of edges to allow arbitrary interleaving of guards and clock changes (reset/release). Our tool, along with the benchmarks used in this paper, is available and can be downloaded from <https://github.com/EQuaVe/GTAReach>. We present selected results in Table 1, with further details in [2].

First, we consider timed automata models from standard benchmarks [21, 34, 39]. Despite the overhead induced by our framework (e.g., maintaining general programs on transitions), we are only slightly worse off wrt. running

time than the standard algorithm, while visiting and storing the same number of nodes. We illustrate this in rows 1–3 of Table 1 by providing a comparison of our tool with the implementation of the state-of-the-art zone-based reachability algorithm using \mathcal{G} -simulation introduced in [24–26].

Next, we consider models belonging to the class of ECA without diagonal constraints. We remark that ours is the first implementation of a reachability algorithm that can operate on the whole class of ECA directly. We compare against an implementation that first translates the ECA into a timed automaton using the translation proposed in [10], and then runs the state-of-the-art reachability algorithm of [24–26] on this timed automaton. From rows 4–7 of Table 1, we observe significant improvements, both in terms of running time as well as number of visited nodes and stored nodes w.r.t. the standard approach.

Finally, in Rows 8–12, we consider the unified model GTA. As already pointed out, model-checking an event-clock specification φ over a timed automaton model \mathcal{A} can be reduced to the reachability on the product of the TA \mathcal{A} and the ECA representing $\neg\varphi$. In this spirit, our implementation allows the model to use any combination of *normal* clocks, *history* clocks, *prophecy* clocks or *timers* and moreover, permits diagonal guards between any of these clocks. To the best of our knowledge, no existing tool allows all these features. We emphasize this by the – in the \mathcal{G} -Sim column of Table 1.

We model simple but useful properties using event-clocks, and check these properties on some standard models from literature such as CSMACD [39], Fire-alarm [35] and Alternating-bit-protocol(ABP) [33]. Note that for the benchmark Fire-alarm-pattern, the specification is modelled using an ECA with diagonals. As a consequence, the product automaton that we check reachability on contains normal clocks and event-clocks. Here, we consider the following ECA specification: no three a 's occur within k time units. The negation of this property can be easily modeled by an ECA with two states and a transition on a with the diagonal constraint $\overleftarrow{a} - \overrightarrow{a} \leq k$, where \overleftarrow{a} is the history clock recording time since the previous occurrence of a , and \overrightarrow{a} is a future clock predicting the time to the next a occurrence. When reading an a , the quantity $\overleftarrow{a} - \overrightarrow{a}$ gives the distance between the next and the previous occurrence. This language is used in [19] to observe that ECA with diagonals are more expressive than ECA. Finally, we remark that the model of ABP contains timers. A more detailed discussion of the model and specifications in these benchmarks is provided in [2].

In conclusion, as can be seen from the experimental results in Table 1, we are able to demonstrate the full power of our reachability algorithm for the unified model of generalized timed automata.

9 Conclusion

The success of timed automata verification can safely be attributed to the advances in the zone-based technology over the last three decades. In fact, [22], the precursor to the seminal works [8, 9], already laid the foundations for zones by describing the Difference-Bounds-Matrices (DBM) data structure. Our goal

in this work has been to unify timing features defined in different timed models, while at the same time retain the ability to use efficient state-of-the-art algorithms for reachability. To do so, we have equipped the model with two kinds of clocks, history and future, and modified the transitions to contain a program that alternates between a guard and a change to the variables. For the algorithmic part, we have adapted the \mathcal{G} -simulation framework to this powerful model. The main challenge was to show finiteness of the simulation in this extended setting. To aid the practical use of this generic model, we have developed a prototype implementation that can answer reachability for GTA. We remark that decidability for GTA comes via zones, and not through regions. In fact, since we generalize event-clock automata, we do not have a finite region equivalence for GTA [28].

We conclude with some interesting avenues for future work. An immediate future work is to use generalized timed automata for model-checking timed specifications over real-time systems. Further, the complexity and expressivity of safe GTA are natural interesting theoretical open questions, but we believe they are not obvious. Both these questions are answered in the timed automata literature using regions. However, we cannot have a region equivalence for our model, since even for the subclass of ECA, it was shown that no finite bisimulation is possible. In particular, it would be interesting to investigate if is possible to have a translation from safe GTA to timed automata. Note that even if such a translation exists, it is likely to incur an exponential blowup since even the translation from ECA to TA costs an exponential. Coming to the complexity of the reachability problem for safe GTA, it is easy to see that our procedure runs in EXPSpace, as we have shown that each reachable zone is a union of equivalence classes of a finite index (see Lemma 13). On the other hand, PSPACE-hardness is inherited from timed automata [6, 8]. Closing the complexity gap is open. We note that even in timed automata, the precise complexity of the simulation based reachability algorithm is difficult to analyze, but its selling point is that it works well in practice. Finally, we would also like to investigate liveness verification for GTA, in particular what future clocks bring us when we consider the setting of ω -words.

References

1. Akshay, S., Bollig, B., Gastin, P.: Event clock message passing automata: a logical characterization and an emptiness checking algorithm. *Formal Methods Syst. Des.* **42**(3), 262–300 (2013)
2. Akshay, S., Gastin, P., Govind, R., Joshi, A.R., Srivathsan, B.: A unified model for real-time systems: Symbolic techniques and implementation. *CoRR abs/2305.17824* (2023)
3. Akshay, S., Gastin, P., Govind, R., Srivathsan, B.: Simulations for event-clock automata. In: *CONCUR. LIPIcs*, vol. 243, pp. 13:1–13:18 (2022)
4. Akshay, S., Gastin, P., Govind, R., Srivathsan, B.: Simulations for event-clock automata. *CoRR abs/2207.02633* (2022)

5. Akshay, S., Gastin, P., Prakash, K.R.: Fast zone-based algorithms for reachability in pushdown timed automata. In: Silva, A., Leino, K.R.M. (eds.) CAV 2021. LNCS, vol. 12759, pp. 619–642. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-81685-8_30
6. Alur, R.: Techniques for automatic verification of real-time systems. Ph.D. thesis, Stanford University (1991)
7. Alur, R., Courcoubetis, C., Henzinger, T.A., Ho, P.: Hybrid automata: An algorithmic approach to the specification and verification of hybrid systems. In: Hybrid Systems, pp. 209–229 (1992)
8. Alur, R., Dill, D.: Automata for modeling real-time systems. In: Paterson, M.S. (ed.) ICALP 1990. LNCS, vol. 443, pp. 322–335. Springer, Heidelberg (1990). <https://doi.org/10.1007/BFb0032042>
9. Alur, R., Dill, D.L.: A theory of timed automata. *Theoret. Comput. Sci.* **126**, 183–235 (1994)
10. Alur, R., Fix, L., Henzinger, T.A.: Event-clock automata: a determinizable class of timed automata. *Theor. Comput. Sci.* **211**(1–2), 253–273 (1999)
11. de Bakker, J.W., Huizing, C., de Roever, W.P., Rozenberg, G.: Real-Time: Theory in Practice: REX Workshop, Mook, The Netherlands. Proceedings, vol. 600 (1992)
12. Bengtsson, J., Yi, W.: Timed automata: semantics, algorithms and tools. In: Desel, J., Reisig, W., Rozenberg, G. (eds.) ACPN 2003. LNCS, vol. 3098, pp. 87–124. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-27755-2_3
13. Bernstein, A.J., Jr., P.K.H.: Proving real-time properties of programs with temporal logic. In: SOSF, pp. 1–11. ACM (1981)
14. Bouyer, P.: Forward analysis of updatable timed automata. *Formal Methods Syst. Des.* **24**(3), 281–320 (2004)
15. Bouyer, P., Chevalier, F.: On conciseness of extensions of timed automata. *J. Autom. Lang. Comb.* **10**(4), 393–405 (2005)
16. Bouyer, P., Colange, M., Markey, N.: Symbolic optimal reachability in weighted timed automata. In: Chaudhuri, S., Farzan, A. (eds.) CAV 2016. LNCS, vol. 9779, pp. 513–530. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-41528-4_28
17. Bouyer, P., Dufourd, C., Fleury, E., Petit, A.: Updatable timed automata. *Theor. Comput. Sci.* **321**(2–3), 291–345 (2004)
18. Bouyer, P., Gastin, P., Herbreteau, F., Sankur, O., Srivathsan, B.: Zone-based verification of timed automata: Extrapolations, simulations and what next? In: FORMATS. LNCS, vol. 13465, pp. 16–42. Springer (2022). https://doi.org/10.1007/978-3-031-15839-1_2
19. Bozzelli, L., Montanari, A., Peron, A.: Taming the complexity of timeline-based planning over dense temporal domains. In: FSTTCS. LIPIcs, vol. 150, pp. 34:1–34:14 (2019)
20. Bozzelli, L., Montanari, A., Peron, A.: Complexity issues for timeline-based planning over dense time under future and minimal semantics. *Theor. Comput. Sci.* **901**, 87–113 (2022)
21. Daws, C., Olivero, A., Tripakis, S., Yovine, S.: The tool KRONOS. In: Alur, R., Henzinger, T.A., Sontag, E.D. (eds.) HS 1995. LNCS, vol. 1066, pp. 208–219. Springer, Heidelberg (1996). <https://doi.org/10.1007/BFb0020947>
22. Dill, D.L.: Timing assumptions and verification of finite-state concurrent systems. In: Sifakis, J. (ed.) CAV 1989. LNCS, vol. 407, pp. 197–212. Springer, Heidelberg (1990). https://doi.org/10.1007/3-540-52148-8_17

23. D'Souza, D., Tabareau, N.: On timed automata with input-determined guards. In: Lakhnech, Y., Yovine, S. (eds.) FORMATS/FTRTFT -2004. LNCS, vol. 3253, pp. 68–83. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-30206-3_7
24. Gastin, P., Mukherjee, S., Srivathsan, B.: Reachability in timed automata with diagonal constraints. In: CONCUR. LIPIcs, vol. 118, pp. 28:1–28:17 (2018)
25. Gastin, P., Mukherjee, S., Srivathsan, B.: Fast algorithms for handling diagonal constraints in timed automata. In: Dillig, I., Tasiran, S. (eds.) CAV 2019. LNCS, vol. 11561, pp. 41–59. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-25540-4_3
26. Gastin, P., Mukherjee, S., Srivathsan, B.: Reachability for updatable timed automata made faster and more effective. In: FSTTCS. LIPIcs, vol. 182, pp. 47:1–47:17 (2020)
27. Geeraerts, G., Raskin, J.-F., Sznajder, N.: Event clock automata: from theory to practice. In: Fahrenberg, U., Tripakis, S. (eds.) FORMATS 2011. LNCS, vol. 6919, pp. 209–224. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-24310-3_15
28. Geeraerts, G., Raskin, J.-F., Sznajder, N.: On regions and zones for event-clock automata. *Formal Methods Syst Design* **45**(3), 330–380 (2014). <https://doi.org/10.1007/s10703-014-0212-1>
29. Herbreteau, F., Point, G.: TChecker. <https://github.com/fredher/tchecker> (v02 - April 2019)
30. ITU-TS Recommendation Z.120: Message Sequence Chart (MSC '99) (1999)
31. Jonsson, B., Vaandrager, F.: Learning mealy machines with timers. Tech. rep. (2018). <https://sws.cs.ru.nl/publications/papers/fvaan/MMT/>
32. Koymans, R., Vytupil, J., de Roever, W.P.: Real-time programming and asynchronous message passing. In: PODC, pp. 187–197. ACM (1983)
33. Kurose, J.F., Ross, K.W.: Computer networking - a top-down approach featuring the internet. Addison-Wesley-Longman (2001)
34. Lugiez, D., Niebert, P., Zennou, S.: A partial order semantics approach to the clock explosion problem of timed automata. *Theor. Comput. Sci.* **345**(1), 27–59 (2005)
35. Muñoz, M., Westphal, B., Podelski, A.: Timed automata with disjoint activity. In: Jurdiński, M., Ničković, D. (eds.) FORMATS 2012. LNCS, vol. 7595, pp. 188–203. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33365-1_14
36. Raskin, J., Schobbens, P.: The logic of event clocks - decidability, complexity and expressiveness. *J. Autom. Lang. Comb.* **4**(3), 247–282 (1999)
37. Sorea, M.: Tempo: A model checker for event-recording automata. Tech. rep., In: Proceedings of RT-Tools'01 (2001)
38. Srivathsan, B.: Reachability in timed automata. *ACM SIGLOG News* **9**(3), 6–28 (2022)
39. Tripakis, S., Yovine, S.: Analysis of timed systems using time-abstracting bisimulations. *Formal Methods Syst. Des.* **18**(1), 25–68 (2001)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

